RFI Response

Response to RFI

NNC05ZVI011L

NASA / Glenn Research Center
Office of Procurement, Cleveland,OH
Global Aeronautical Network Requirements

Specific to Telecommunications Mobility for IPSEC Environments

NASA
Glenn Research Center

Deltao, Inc.
High-end Network Design and Optimization
www.deltao.com
GSA Schedule 70
GS-35F-0378R

CCIE No. 5186

Prepared by:
Vince Fortunato
Sr. Network Engineer
vfortunato@deltao.com
703-589-6298

Last Modified March 22, 2005
Version 1.0

I. Overview

The goal of this response is to bring to light the specific network protocols that can coexist with various lower-layered transport mediums inherent to the Global Airspace System. These technologies and protocols exist today as separate entities. However, the requirements within the associated RFI indicate vast complexities which force these entities to overlap. From the standpoint of allowing layer 3 transport technologies through this system, a high-end network designer that is intimately familiar with virtually all of the modern network disciplines is required. Some of these disciplines include (just to name a few):

- Encapsulation of generic tunnels into IPSEC
- SA (Security Association) behavior when approaching various network-to-network limits
- Fragmentation issues when dealing with legacy applications, MTU's, an TCP maximum segment sizes
- Flexible build/tear-down tunnel requirements, with quick routing convergence

As you can see, the Global Airspace System must meet requirements pertaining to flexibility, security, and interoperability. Typically, IP encryption mechanisms such as IPSEC do not allow for the flexibility required when connecting mobile entities. By itself, IPSEC has stringent endpoint configuration and does not allow for multicasting. However, with slight variations to these technologies, the flexibility requirements can be met.

II. Design Abstract

Obtaining Network Flexibility: GRE Tunneling

A generic IP tunneling mechanism must be used as the first "envelope" for other types of encryption. GRE (Generic Routing Encapsulation) can accommodate multicasting. Therefore, GRE tunnels can be used to carry routing protocols for IP mobility, and DLSW (Data Link Switching) traffic for time-sensitive, low bandwidth protocols such as Air Traffic Management (ATM). GRE tunnels can more easily accommodate failover and high-availability scenarios via multiple floating static routes. GRE tunnels would give the Global Airspace System the required flexibility without sacrificing worldwide standards-based networking protocols.

Modern routing protocols (such as OSPF and EIGRP) require layer 3 multicasting to work properly. By its very nature IPSEC does not allow broadcast or multicast traffic through the network. For this reason, GRE (Generic Routing Encapsulation) Tunnels should be used. These GRE tunnels will be encapsulated inside IPSEC tunnels. This encapsulation will accomplish numerous simultaneous functions. Some of these functions include:

- Maximum of three SA's (Security Associations) per IPSEC connection (reduces processor utilization)
- Ability to send multicast and broadcast traffic (and hence OSPF) through IPSEC
- Ease of routing manipulation/floating static routes for failover configuration

Obtaining Network Security: IPSEC VPN's

VPN technology has enjoyed enormous gains in the remote access world. However, VPN as a pinned-up enterprise solution can produce many pitfalls:

- Poor performance ("not enough bandwidth" perception)
- Routing protocols could not easily traverse a hub-and-spoke network, especially spoke-to-spoke traffic
- Configuration of the VPN encryption and associated encapsulated tunnels was complex and not-scalable

After digging deeper into the issues, it is apparent that IPSEC technology is not normally understood to the depth required to "optimize" a VPN tunnel for a particular application. Usually bandwidth gets the blame for poor performance. The fact is, bandwidth is not the problem. The problem is derived from a combination of *router CPU utilization* and *packet fragmentation*. Traditional routers have used their CPU for regular routing functions as well as encryption and decryption duties. The latter can bring the CPU to 99% very quickly. If sustained for several seconds, the buffers will begin to overflow, and the tunnel drops. When the tunnel drops, default timers can also lead to rather lengthy recovery periods. As TCP retransmits, we can get a back-log of traffic, and consequently a very slow response time which ultimately falls off the user acceptance charts. The amount of packet encryption is also directly related to the packet sizes and variance of packet sizes (a flow of mixed-sized traffic such as 86 bytes, 1500 bytes, and 300 bytes will produce an enormous difference in performance). Couple the packet-sizing issues with fragmentation issues (where large 1500-byte packets can't fit into the tunnels with smaller MTU sizes), and performance problems may be compounded, especially if the application has the DF bit ("don't fragment" bit) set. Thus, you can see how even protocols (such as

Air Traffic Management) are susceptible to poor performance over IPSEC when the tunnel MTU requirements and buffer allocations are not optimized correctly.

Putting everything together: DMPVPN (Dynamic Multipoint VPN's)

The usual drawback with GRE tunnels encapsulated inside of IPSEC tunnels is the complexity of the configuration. Thus, the solution does not scale when new endpoints must be added and/or dropped (a desired trait with mobile entities). Normally, a GRE/IPSEC tunnel is configured with known tunnel endpoints. This drawback fails the specific Global Airspace System design requirement of using host locators unbound from IP addresses. However, we can overcome this drawback of GRE tunneling by using NHRP (Next Hop Routing Protocol). Basically, NHRP allows the distant end IP address of any tunnel to be variable and/or dynamic. IPSEC encryption SA's can then bind to a certificate authority. This scenario will work as long as there is a hub-and-spoke situation. In effect, this will always work because we can ensure the mobile entity (aircraft) is the hub of a new IPSEC network. Spoke-to-spoke (ground to ground) routing can also be accommodated with on-demand IPSEC links and routing. Each mobile entity will thus open up and tear down entire hub-and-spoke communications networks at will. Below is a general depiction of a possible network design scenario.

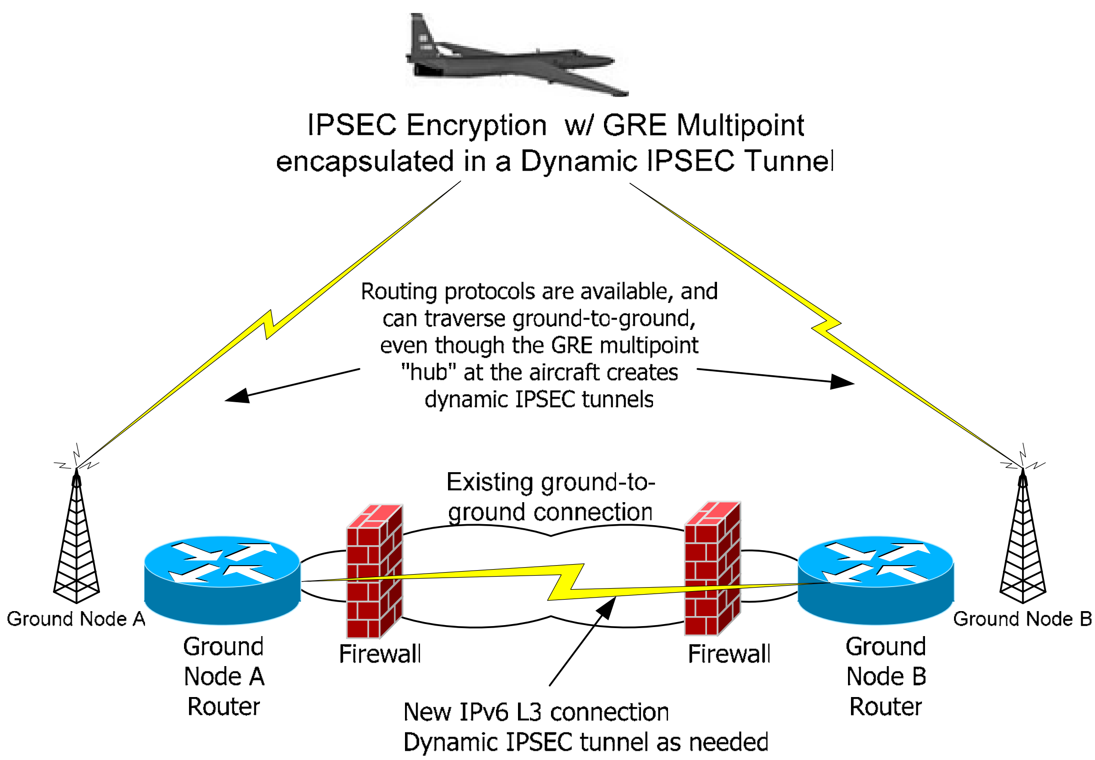

IPSEC configuration involves many levels of configuration, and is tightly configured on each and every router. The “spoke” or “ground” routers’ IP addresses will be quite various for the numerous global networks. The answer to this is a new technology called DMVPN (Dynamic Multipoint VPN) tunneling. A VPN router can be bootstrapped and/or preconfigured at a central location. When connected at a remote site, NHRP (Next Hop Routing Protocol) is used to learn the IP address of that site. Once the IP address is learned, IPSEC tunnel source and destination IP addresses are automatically updated at the hub location, and GRE access-lists are automatically configured to match on IPSEC traffic. Keys are then exchanged, and IPSEC tunnels are established. The new networks are automatically entered into the routing protocol database. Spoke routers with direct connections will also be automatically configured in a similar fashion once the main tunnels to the hub router(s) are up. The shortfalls of OSPF with NBMA (Non Broadcast Multi Access) networks (Frame Relay and ATM) are also recognized. The complete design specifics can be made with further discussion.

III. About Deltao

Deltao is currently working a detailed white paper on the complexities and interdependencies of Dynamic Multipoint VPN tunneling for a major enterprise organization in support of a project mandated by the Department of Homeland Security. We are a veteran-owned small business and operate under GSA Schedule 70, SIN 132-51, Contract No. GS-35F-0378R. The founder, Vince Fortunato, holds a degree in Engineering Physics from West Point as well as the networking industry’s most acclaimed certification: the CCIE (Cisco Certified Internetwork Expert). Vince’s background involves over a decade of experience on several national projects ranging from government to large enterprise. He is a former Army Signal Officer, and was involved in several tactical communication projects at Fort Huachuca, Arizona. Deltao’s niche market is high-end network design. Our capabilities include core network redesigns and optimization, as well as VoIP, VPN, and PIX firewalls. We typically perform initial network assessments, and move directly into design and implementation. On some complex designs, we will include a testing or a prototype phase. We bring extreme versatility and flexibility to a firm’s overall network capability. Deltao also brings industry commercial best practices to government facilities. Deltao can perform the telecommunication assessment, design, and implementation of the GAS.